A Sink-driven Approach to Detecting Exposed Component Vulnerabilities in Android Apps

Daoyuan Wu, Xiapu Luo, and Rocky K. C. Chang
Department of Computing
The Hong Kong Polytechnic University
{csdwu, csxluo, csrchang}@comp.polyu.edu.hk

arXiv:1405.6282v1 [cs.CR] 24 May 2014

ABSTRACT

Android apps could expose their components for cooperating with other apps. This convenience, however, makes apps susceptible to the *exposed component vulnerability* (ECV), in which a dangerous API (commonly known as *sink*) inside its component can be triggered by other (malicious) apps. In the prior works, detecting these ECVs use a set of sinks pertaining to the ECVs under detection. In this paper, we argue that a more comprehensive and effective approach should start by a systematic selection and classification of *vulnerability-specific sinks* (VSinks). The set of VSinks is much larger than those used in the previous works. Based on these VSinks, our sink-driven approach can detect different kinds of ECVs in an app in two steps. First, VSinks and their categories are identified through a typical forward reachability analysis. Second, based on each VSink's category, a corresponding detection method is used to identify the ECV via a customized backward dataflow analysis. We also design a semi-auto guided analysis and validation capability for system-only broadcast checking to remove some false positives. We implement our sink-driven approach in a tool called ECVDetector and evaluate it with the top 1K Android apps. Using ECVDetector we successfully identify a total of 49 vulnerable apps across all four ECV categories we have defined. To our knowledge, most of them are previously undisclosed, such as the very popular Go SMS Pro and Clean Master. Moreover, the performance of ECVDetector is high, requiring only 9.257 seconds on average to process each component.

1. INTRODUCTION

To ease and accelerate the development of apps, Android takes a modular programming paradigm that empowers developers to focus on essential building blocks (i.e., components [1] in Android terminology). Moreover, Android apps could expose their components for cooperating with other apps. For example, both Facebook and Twitter apps leverage an existing photo-capturing component exposed by a camera app by simply sending a request to it. The photo-capturing component in this case is an *exposed component* that serves external requests from other apps.

In return for this convenience, exposed components, if not well designed, might run into security risks. In fact, vulnerabilities might exist if a dangerous API inside exposed components can be triggered by other (malicious) apps. We refer to this class of vulnerabilities as *exposed component vulnerability* (ECV), and the dangerous APIs in ECVs are the sinks of potential attack flows. Usually, ECVs could be exploited by an attacker to perform dangerous operations by simply sending crafted inputs from a regular app to a victim app, both installed on the same phone.

Several methods for detecting specific ECVs [19, 7, 26, 36, 32, 33] have been proposed in the past. In all of these works, detecting these ECVs use a set of sinks pertaining to the ECVs under detection. Specifically, Woodpecker [19], DroidChecker [7], and the very recent IntentFuzzer [33] are designed to detect *permission leakage* in Android apps, and they focus on a specific kind of sinks that would directly leak permissions once victim apps are exploited. A recent work, SEFA [32], further considers some database-related sinks that are aimed by ContentScope [36] for a special kind of components (i.e., Content Provider). Finally, CHEX [26] discovers potential vulnerabilities related to another kind of sinks, which are the data sinks that might cause unauthorized read or write operations on sensitive resources.

In this paper, we argue that a more comprehensive and effective approach should start by a systematic selection and classification of sinks. Note that in the context of ECV detection, sinks should be vulnerability-specific (i.e., vulnerability-specific sinks, or VSinks) in contrast to the general data sinks for privacy leak detection [2]. This approach will help resolve two major issues in the previous detection methods. First, the set of sinks obtained from our approach is much more comprehensive. It will therefore help the previous methods to discover new ECVs. Another and also more important issue is that the prior methods are tightly coupled with individual analysis requirements of their selected sinks. They therefore cannot collaborate with one another to form a more general detection method. Our approach, on the other hand, breaks this coupling by admitting different kinds of sinks and categorizing them for different analysis methods.

Using this sink-driven approach, we adopt a systematic strategy to select VSinks and classify them into multiple categories according to their different analysis requirements. In this strategy, we combine multiple metrics (e.g., permission

semantics and API names) to systematically define rules. These rules are made according to a simple, but practical, rule syntax. We further write a rule interpreter to automatically select and classify VSinks according to the defined rules.

Based on the categorized VSinks, our sink-driven approach can detect different kinds of ECVs in an app in two steps. First, VSinks and their categories are identified through a typical forward reachability analysis. We employ an iterative intra-procedural algorithm with flow sensitivity to perform this reachability analysis. Second, based on each VSink's category, a corresponding detection method is used to identify the ECV via a customized backward dataflow analysis. The backward, instead of prior forward, dataflow analysis is chosen to adapt more categories of sinks and organically cooperate with the forward analysis. Furthermore, we design a semi-auto guided analysis and validation capability for system-only broadcast checking for removing some false positives.

In summary, this paper makes the following contributions:

- **Methodology.** We propose a new sink-driven approach to systematically tackling the ECV detection problem (§3). This approach includes a systematic strategy for VSink selection and classification, a general detection method to identify all categories of potential ECVs, and semi-auto guided analysis for excluding some sink-specific false positives.
- **Tool and dataset.** We implement our sink-driven approach in a tool called ECVDetector (§4). We also design three analysis enhancements in ECVDetector, and the major one is that ECVDetector can validate broadcast checking, a capability that could significantly reduce false positives. Moreover, ECVDetector identifies a total of 372 VSinks across four categories, as well as 183 data source APIs. We are going to release this dataset to the Android research community.
- **Evaluation and results.** We evaluate ECVDetector with the top 1K Android apps from Google Play (§5). In total, we identify 49 vulnerable apps across all the four ECV categories. To our knowledge, most of them are previously undisclosed, such as the very popular Go SMS Pro and Clean Master. Moreover, the performance of ECVDetector is high, requiring only 9.257 seconds on average to process each component.

2. BACKGROUND

2.1 Exposed Component

Exposed components are a subset of Android-defined *exported components* (i.e., components that other apps could access). Some exported components have reliable permissions to protect them, but the exposed components in our threat model are fully exposed to other zero-permission apps or apps with only normal level permissions. This threat model is also adopted in previous related work [19, 36]. We detect exposed components according to the following rule:

RULE 1 (exposed component determination). A component is considered as an exposed component when it satisfies both two conditions:

C1: It must be enabled so that it could be successfully instantiated by the system; AND

C2: It must be explicitly or implicitly exported so that other app could access it without permission or only with `normal` *level permission.*

Each Android app contains an `AndroidManifest.xml` file, which defines a set of component attributes. Therefore, the rule for the exposed component determination is to inspect corresponding attributes. Using *C1*, we can exclude those useless components (i.e., those who set `enabled` attribute as false), such as the already deprecated components. For *C2*, the *explicitly exported component* are those with their `exported` attribute set to true. The *implicitly exported components* are by default exported by Android convention. For example, Intent-based components with `intent-filter` tag and Content Provider components without `exported` attribute are implicitly exported (Note that this default Provider convention is disabled since Android 4.2, but all other lower Android versions have to be compatible with this convention.).

2.2 Problem Statement

Exposed Component Vulnerability, ECV, is one kind of the classic confused deputy vulnerabilities [21]. The exposed component mechanism in Android allows other apps to send a request or input to the victim component. However, if the victim component cannot differentiate whether the request is from trusted parties or not and blindly execute its own code for finishing the request, then it becomes as a confused deputy. Furthermore, if the triggered "deputy" would execute some security-related operations, then an ECV exists. In this case, the victim component is a stepping stone for attackers and also the direct executor of attack behaviors. A model of such attack scenario could be found in [15]. Essentially, the ECV is originated from insecure IPC communication in Android (mainly delivered by *Intent* this IPC object), and a related problem is the attack on unauthorized Intent receipt [8]. However, this issue is mainly due to the ambiguity during the resolution of Intent messages, thus not related to the ECV.

The impact of ECVs largely depends on the triggered sinks and the working mechanism of victim app itself. Some ECVs are serious (such as forcing the victim app to send a SMS), while others could be even treated as bugs (e.g., starting a private Activity for attackers). It is nearly impossible to give a clear boundary between such vulnerabilities and bugs. In this paper, we generally consider these two kinds as vulnerabilities: (1) those cause security or privacy risks to the user or phone, and (2) those have security impact to internal status of victim apps. Moreover, issues in exposed Activity components that require user interaction are not included in our threat model, since they are not stealthy and usually not recognized by vendors.

Our scope of ECV detection in this paper is the same as two previous related work [19, 26] that use the sink-based flow analysis. However, only detecting attack surfaces and giving warnings (one focus of [8, 28]) is not enough for our goal. Instead, in this work we try to discover the real exposed component vulnerabilities. On the other hand, the complicated attacks on unauthorized origin crossing [31] is out of our scope, because it relies on too much browser-specific domain knowledge. Moreover, two technical issues (Java reflection and native code), are out of our paper scope. This also implies that we only need to select sinks from documented APIs (i.e., those could be found in Android SDK).

2.3 VSink and its Taxonomy

VSink, short for Vulnerability-specific Sink, is a sink API where an attack flow could go to and related to ECVs. In this paper, we consider VSinks mainly from these two channels:

C1: APIs that will cause security risks to permission-protected or privileged resources; OR

C2: APIs that will make security impact on internal status of apps.

VSinks can be further classified into four categories according to their different analysis requirements. We introduce these four categories as follows.

`VS_Direct.` This category of VSinks is usually related to privileged resources, and they can be directly used to launch an attack. For instance, the `removeAccount` is a `VS_Direct` sink. Another example is `SmsManager.sendTextMessage()`, which can send a SMS message to outside without the attention of user. Analyzing `VS_Direct` sinks is relatively straightforward, usually based on a path reachability analysis. A major difference between VSinks and sinks in other Android works is that not all permission-required APIs will be considered as VSinks, only the vulnerability-related ones. For example, `WAKE_LOCK` and `VIBRATE` APIs are excluded.

`VS_DirectByParam.` This category of VSinks is similar to `VS_Direct`, but they rely on the incoming parameters to exhibit different attack behaviors. Different parameters could cause different attack consequences. For instance, `ContentResolver.delete(Uri)` could exhibit different attack behaviors with different URI parameters, such as "content://sms" for deleting SMS. Analyzing such ECVs is more complex than `VS_Direct`, because a chain of variable dependences must be investigated, so that we can obtain the value of parameters. Note that this category of VSinks has not *generally* been considered in prior works (e.g., CHEX [26] only consider some similar APIs when external Intent can control their parameters, while such relationship between Intent and parameters in our `VS_DirectByParam` is not necessarily required).

`VS_Input.` This category of VSinks mainly fulfills the goal of misusing privileged resources, and this kind of misuse relies on attack inputs to flow into `VS_Input` sinks. Network-related sinks are the typical examples of `VS_Input`, such as `HttpClient.execute()`. Once attack inputs flow into them, they could be exploited to misuse protected Internet resources. Besides network-related `VS_Input`, APIs, such as `startService()` and `startActivity()`, also belong to this category. Note that although `VS_Input` sinks usually contain parameters, this is not necessarily required. For instance, inputs flow into the caller object of `HttpURLConnection.connect()` will also cause `VS_Input` ECVs.

`VS_Public.` VSink APIs in this category are those which might not directly transmit privileged resources to attackers but will make them public to other apps. An attacker could then leverage a local app to steal privileged resources outputted from these `VS_Public` APIs. One typical kind of `VS_Public` sinks is the output methods defined in the `android.util.Log` class, such as debug-level method `d(String tag, String msg)`. Vulnerable components might put privileged resources (e.g., GPS locations) to these output methods, and an attacker could collect these log outputs in runtime.

Besides VSinks, data source APIs are also related to specific ECV detection (e.g., `VS_Public` ECVs). In this paper, we consider the data source APIs that can read permission-protected resources and call them `VSource`, such as the `BLUETOOTH` APIs.

3. ECVDetector DESIGN

In this section, we present our design of ECVDetector, which implements the sink-driven approach to systematically tackling the ECV detection problem. As shown in Figure 1, ECVDetector first selects and classifies VSinks. Then based on these VSinks, we propose a general detection method to identify all categories of potential ECVs. This method organically combines forward reachability and backward dataflow analysis, and drives them by the characterized VSinks. In particular, we take the backward, instead of previous forward, dataflow analysis for adapting more categories of sinks, such as the `VS_DirectByParam` category.

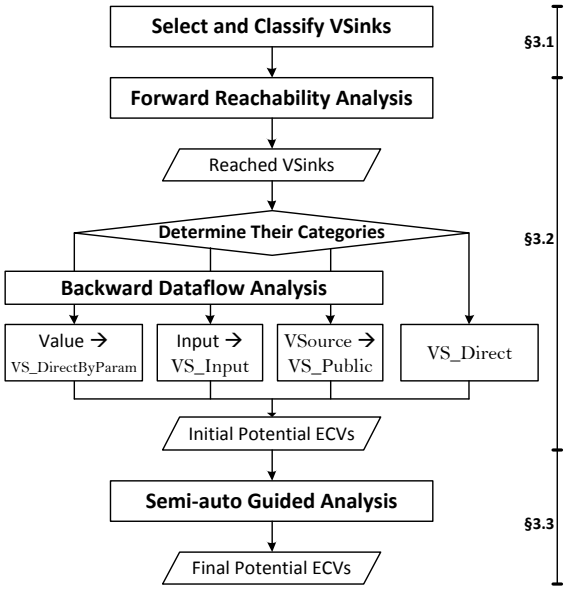

Figure 1: The overall work flow of ECVDetector.

More specifically, the first step produces a number of categorized VSinks, which are the dataset for the subsequent analysis. Then ECVDetector employs a core module with a typical forward reachability analysis to tackle the common analysis task (i.e., identifying all reachable VSink calls from attack entry points), and further leverages a customized backward dataflow analysis to handle different categories of VSinks in several dedicated modules. As the last step, a semi-auto guided analysis is conducted for excluding some sink-specific false positives. In the following three sections, we will discuss each critical analysis phase in more detail.

3.1 VSink Selection and Classification

We propose a systematic strategy for VSink selection and classification. The basic idea of this strategy is that we combine multiple metrics (e.g., permission semantics and API names) to systematically define rules. These rules are made according to a simple, but practical, rule syntax. We further write a rule interpreter to automatically select and classify VSinks according to the defined rules.

Most of VSinks are selected from permission-protected or privileged APIs, according to the *C1* channel in §2.3. However, Android does not provide a complete and accurate mapping for permissions and their corresponding API calls. In fact, permission information provided by Android developer document is limited and might contain errors [14]. However, two great works [14, 3] have attempted to address this limitation. Specifically, Stowaway [14] constructs mappings for Android 2.2 framework using dynamic API fuzzing. In contrast, PScout [3] employs a version-independent static analysis method to extract permission specifications from multiple versions of Android.

In this paper, we adopt API-to-permission mappings from Stowaway instead of PScout for two reasons. First, although PScout provides a significant number of undocumented APIs, we only select VSinks from the range of documented APIs (see §2.2). In terms of documented APIs, PScout does not provide more permission mappings than Stowaway. Second, we find Stowaway produces a more accurate mapping than PScout, in terms of fewer false positives. This might due to the fact that nearly every mapping found by Stowaway is verified by dynamic execution. The accuracy of permission mapping is important for our VSink selection, because any incorrect VSink will directly introduce false positives to our ECV detection. Therefore, currently we only adopt the mappings from Stowaway. In the future, we could also take advantage of the version-independent feature of PScout.

We obtain a total of 456 validated documented APIs from Stowaway. Assessing each API semantic is practically infeasible for two reasons: (1) it would take a quite large manual workload, and (2) two much manual analysis without a principle may incur some inaccuracy (e.g., assign wrong VSink tags). Therefore, we propose to combine multiple metrics for systematic selection and classification. These metrics include permission semantics and levels, API names, parameter and return-value types. Among all metrics, permission semantic is the major metric for our system. In fact, we only extract 56 kinds of permissions from those 456 documented APIs. Therefore, we can divide them into 56 clusters (if an API needs multiple permissions, we choose its first marked permission), because APIs marked with the same permission are likely to share similar VSink nature. For instance, combining with quick scan of API names, we find all 6 APIs under `SEND_SMS` permission can be directly categorized into `VS_Direct`.

To facilitate multiple metrics based selection and classification, we further design a simple, but practical, rule syntax, as illustrated in Figure 2. Each rule is a five-element tuple, which describes what tag a particular API would be assigned, when one metric of this API satisfies a special string pattern. We define four kinds of pattern-matching actions, such as "START" for "start with" action. Finally, the tag in the syntax mainly includes four VSink categories. Besides them, we also define a special tag named "Tag_Delete", which can be used when we want to exclude an API. Based on this syntax, we define four general rules and a set of permission-specific rules for extracting categorized VSinks from permission mappings. We then write a rule interpreter to automatically select and classify VSinks according to the defined rules.

Besides the privileged APIs, another channel to select VSinks is the APIs that would make security impact on internal status of apps. Three general kinds of such APIs have been considered in our work. The first kind is the file operation APIs, such as the `write` and `append` methods from the `java.io.Writer` class. Another kinds of APIs is the database-related APIs, such as those from `SQLiteDatabase` and `ContentResolver` classes. The third one is the logcat APIs, i.e., logging APIs from the `android.util.Log` class. We then design three additional rules to automatically join them into our VSink set. Moreover, APIs from the `AndroidHttpClient` class (missed by both Stowaway and PScout) are similarly complemented into our final VSink results.

<rule> ::= <metric> <action> Pattern TAG <tag> (1)
<metric> ::= @class | @method | @param | @return (2)
<action> ::= START | IS | CONTAIN | END (3)
<tag> ::= VSink | Tag_Delete (4)

Figure 2: The rule syntax for VSink selection and classification.

3.2 Forward and Backward Analysis

3.2.1 Forward Reachability Analysis

We employ an iterative intra-procedural algorithm with flow sensitivity to perform reachability analysis (summarized in Algorithm 1). Note that since reachability analysis only relies on control flow, context-insensitive analysis is enough in our forward module. The algorithm first constructs a control flow graph (CFG) for each method, and then traverses every statement in a particular order according to the search strategy (e.g., DFS). For each call site, we determine whether it is a VSink or can be resolved into a method defined by the app, respectively. In particular, this call site resolving procedure is performed on demand along the forward analysis. We maintain two lists for caching reached VSinks and resolved methods. After the lists stop growing, we obtain all reachable VSinks and then invoke specific modules to further process them. Finally, the algorithm will handle each resolved method in the same way.

We generate call chains to facilitate inter-procedural backtrack analysis. A call chain is a path of call graph, from an entry caller to an ending callee. Generating individual call chains, instead of a whole call graph, could ease the path track procedure of subsequent modules and allow them concentrate on the design of flow analysis. In contrast, traversal between two nodes on call graph might involve several parallel paths. Moreover, an additional dataflow fact *join* operation needs to be considered in the dataflow analysis using non-linear graph traversal.

Our call chain captures not only caller and callee methods, but also precise calling context information by recording call-site heap locations and their call strings (both are not shown in Algorithm 1, for simplicity). On one hand, the heap location context information is essential for backward modules to jump back to each original call site. Otherwise, ambiguity might arise when a caller method contains multiple similar call sites targeting the same callee. On the other hand, we leverage call string to avoid re-analysis of callee method with the same dataflow value context. More specifically, with the help of SSA IR form[1], we can distinguish call

[1]SSA represents static single assignment, while IR is short for intermediate representation.

Algorithm 1 Forward Reachability Analysis

Input: *entryPts* = [entry point], *vSinks* = [VSink]

```
1: for all ep ∈ entryPts do
2:     IterativeForward(ep, initchain)
3: end for
4:
5: procedure IterativeForward(method, chain)
6:     reachedS = [], nextM = [], chain.add(method)
7:
8:     cfg ← BuildCFG(method)
9:     for all cs ∈ CallSites(DFS(cfg)) do
10:        if cs ∈ vSinks then
11:            reachedS.add(cs)
12:        else
13:            rc ← ResolveCall(cs)
14:            if rc ∈ [app defined method] then
15:                nextM.add(rc)
16:            end if
17:        end if
18:    end for
19:
20:    for all r ∈ reachedS do
21:        category ← JudgeCategory(r, vSinks)
22:        InvokeSpecificModule(category, r, chain)
23:    end for
24:    for all m ∈ nextM do
25:        IterativeForward(m, ShadowCopy(chain))
26:    end for
27: end procedure
```

Algorithm 2 Backward Dataflow Analysis

Input: *cy*: category, *r*: a reached VSink call, *chain*, *vSrc*

```
1: ii ← chain.size()-1, itv ← GetTVarsFromCallSite(r)
2: IterativeBackward(ii, itv, r)

   ▷ i: index, tv: tainted variables, cc: calling context
3: procedure IterativeBackward(i, tv, cc)
4:     sb ← chain.get(i).GetSSABody()
5:     ts ← InitTaintSet(), ts.join(GetTaintedVars(tv))
6:     cs_start ← GetStartCallSite(cc)
7:     ts.join(GetTVarsFromCallSite(cs_start))
8:
9:     cs_old ← cs_start
10:    while sb.hasPreviousCallSite(cs_old) do
11:        cs_new ← sb.getPreviousCallSite(cs_old)
12:        if isThisCallSiteTainted(cs_new, ts) then
13:            ts.join(Determine-PropagateTaint(cs_new,
   ts))
14:            MarkResult_VSrc(cy, cs_new, vSrc)
15:            MarkResult_Constant(cy, cs_new)
16:        end if
17:        cs_old ← cs_new
18:    end while
19:    MarkResult_Input(cy, ts, chain)
20:
21:    if isContinueIterative(i, ts, cy, chain) then
22:        lcc ← chain.get(i).GetLastCallContext()
23:        IterativeBackward(i-1, TransformTVars(ts), lcc)
24:    end if
25: end procedure
```

sites with different entry dataflow values through simply investigating their call strings.

3.2.2 Backward Dataflow Analysis

Except for `VS_Direct`, the other three kinds of VSinks require further backward dataflow analysis. Although each backward analysis is independent for a dedicated task, they share a common backward dataflow analysis method. We thus first design this common backward method and then apply it to the three dedicated modules.

Our backward method relies on call chain (generated by forward module) to achieve inter-procedural context-sensitive backtrack analysis (shown in Algorithm 2). The core of this algorithm is an iterative backward analysis procedure. This iterative procedure starts with extracting SSA-form method body from call chain according to current chain index. Then, it initializes a intra-method taint set and joins the incoming tainted variables. After this, we need to locate the starting call site according to the provided calling context. If calling context is the null (this can happen when we backtrack between two originally disconnected callback functions), we directly take the last call site as our starting point. Similarly, we join the new tainted variables obtained from starting call site into the taint set. We then loop all call sites before the starting point. For each tainted call site, we further determine whether we need to propagate the taint into new variables. Moreover, different dedicated modules may choose to mark result for tainted `VSource` or constants at this time. After the loop, some modules would further inspect whether there are tainted inputs. Finally, the algorithm will judge whether it needs further backtracking, according to the current index number and variables in taint set. For example, if we already backtrack to the first method in chain, or no parameters and fields tainted, the algorithm will terminate.

Generating proper taint objectives is important for our backward analysis. First, we need to obtain appropriate initial taints from reached VSink calls. Since our current VSinks have no fine-grained information to indicate which parameter is critical, we then take a conservative approach that taints all encountered parameters to avoid any false negative. Moreover, some VSink APIs have no parameters involved (e.g., previously mentioned `HttpURLConnection.connect()`), we thus taint their caller object. Second, we need to maintain a mapping for tainted parameters during the procedure of method switching, so that we can obtain the correct variable format under different SSA method bodies.

We further design three dedicated modules for tackling ECV detection problems under `VS_DirectByParam`, `VS_Input`, and `VS_Public`. With the help of proposed backward method, these modules are relatively easy to design. Specifically, for `VS_DirectByParam`, we mark the result from tainted constants and inputs. Because sometimes static analysis cannot obtain accurate parameter values, e.g., when they rely on dynamic execution outputs. Therefore, we adopt every tainted constants into `VS_DirectByParam` result, to mimic the ideal values. While for `VS_Input` and `VS_Public`, our main task is to backtrack whether there are tainted inputs and `VSource`, respectively .

3.2.3 Analysis Enhancements

We also design three kinds of enhancements to the basic forward and backward analysis in ECVDetector. In general,

these enhancements can help ECVDetector reduce false positives, avoid unnecessary analysis overhead, and output more expressive result logs.

The first enhancement is to validate system-only broadcast checking in the flow analysis. Broadcasts are system-wide events (e.g., battery is low), which will be delivered to registered Broadcast Receivers when the corresponding events occur. For example, an Broadcast Receiver with the name `com.example.BootReceiver` (see Figure 8 in the Appendix) registers the `BOOT_COMPLETED` broadcast, which will then trigger `BootReceiver` when the system has booted. Note that although third-party apps can also declare their own broadcasts, many broadcasts are only sent by the operating system and prohibited by non-system senders [14]. Therefore, attackers cannot inject a fake broadcast Intent with `BOOT_COMPLETED` as the action name into `BootReceiver`. However, it is still insufficient to prevent external crafted inputs, since attackers can directly trigger `BootReceiver` by set the explicit Intent target name. A safe and common way to mitigate this issue is to explicitly check the action name of system broadcasts in the code, as suggested in [8] and [28]. A typical code pattern for such system-only broadcast checking is illustrated in Figure 9 (see Appendix).

Since broadcast checking can efficiently protect exposed Broadcast Receivers, we thus design a validation capability for system-only broadcast checking in ECVDetector's flow analysis, as an enhancement to the forward analysis in §3.2.1. Besides helpful to reduce false positives, this validation can also improve the system performance, because ECVDetector can avoid the subsequent analysis once it identifies checking. Our validation is targeted at the code pattern in Figure 9 (i.e., the If-Else checking using `equals` API in Java's `String` class), since it is the most straightforward way to perform broadcast checking. We also believe it is easy for ECVDetector to cover other kinds of checking patterns, once their domain knowledge is provided.

To facilitate accurate validation, a critical job is to collect sufficient system-only broadcasts. Note that our goal is to identify enough broadcasts that could cover most app cases, and proposing a new way to dig out a complete set is out of the paper scope. There are two existing related resources we can leverage. First, the `AndroidManifest.xml` file in each version of Android source code defines a list of system broadcasts with the tag name `protected-broadcast`. We thus collect 133 such broadcasts from the recent Android 4.3 platform. Second, Stowaway [14] uncovers 62 system broadcasts by dynamic testing, including those dynamically declared in the code. We then merge these two broadcast sets into a new one, and finally obtain in total of 143 unique system-only broadcasts for ECVDetector.

Second, we avoid backtracking some uncritical parameters to reduce overhead. A typical example is the first `tag` parameter of logcat APIs, such as `Log.v(String tag, String msg)`. Developers usually assign insensitive values (e.g., string constants) into this parameter, thus there is no need to analyze it. Otherwise, the backward module may waste tracing several methods before arriving its initialization method.

Third, we output input-related variable values for more expressive result logs. Our backward analysis could taint a dependence between the input and parameter, like CHEX [26]. However, such coarse-grained dependence without detailed propagation knowledge sometimes is not enough. To this end, we choose to output some input-related variable values into our result logs, such as the log like `Input:r4 = r2.getStringExtra("referrer")`, where r2 is the original Input or Intent object. In this way, we can obtain more expressive and meaningful result logs.

3.3 Semi-auto Guided Analysis

We need to further filter some false positives in the initial set of potential ECVs identified by previous analysis. These false positives are mainly from two VSink categories: `VS_DirectByParam` and `VS_Input`. The reason is that VSinks in these two categories generally rely on parameters to exhibit their specific behaviors. Therefore, analyzing these VSinks is usually sink-specific and parameter-specific, meaning that we have to combine detailed sinks and their parameter values for detection. However, it is nearly impossible for ECVDetector to automatically handle it, because too much domain knowledge is needed.

To address this issue, we propose semi-auto guided analysis to quickly filter false positives. Its basic process is illustrated in Figure 4. In particular, we take advantage of the result logs outputted by ECVDetector. Figure 3 shows the logs of example false positives in the `VS_DirectByParam` category. Specifically, we first collect all unique results logs from `VS_DirectByParam` and `VS_Input`. Then, we conduct manual analysis to find all false positive logs and their patterns. This step largely relies on expert knowledge. Little effort is actually required because many logs are similar. Finally, we apply our extracted false positive log pattern to all related apps and take scripts to automatically filter those matched cases.

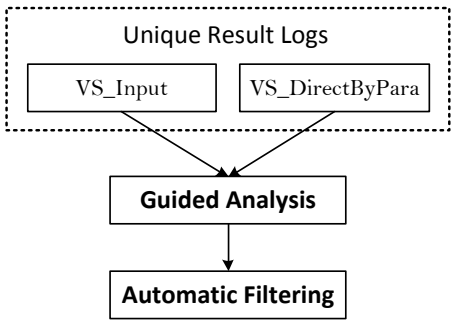

Figure 4: Basic process of semi-auto guided analysis.

4. ECVDetector IMPLEMENTATION

We have implemented ECVDetector with 4,565 lines of Java code, Python scripts and shell scripts. As shown in Figure 5, ECVDetector consists of four components. Specifically, Manifest Analyzer is implemented only in Python scripts, while the other three components (i.e., Entry Point Locator, Vulnerability Analyzer and VSink Selector) are based on the Soot framework [24]. ECVDetector contains four execution steps: one preparation step (i.e., step 0 at the bottom of Figure 5) and three analysis steps. The preparation step is executed only once for generating VSinks. Then, ECVDetector analyzes each Android app for ECVs in three consecutive analysis steps.

VSink Selector. By running VSink Selector with the inputs of Android framework code and Stowaway permission mappings, we obtain a total of 372 categorized VSinks.

"mount" to '$r10 = virtualinvoke $r9.<java.lang.Runtime: java.lang.Process exec(java.lang.String)>("mount")'
"logcat -d " to '$r14 = virtualinvoke $r3.<java.lang.Runtime: java.lang.Process exec(java.lang.String)>($r13)'
"referrer" to 'virtualinvoke r0.<android.database.sqlite.SQLiteDatabase: int delete(...)>("referrer", null, null)'

Figure 3: Logs showing example false positives in `VS_DirectByParam` category.

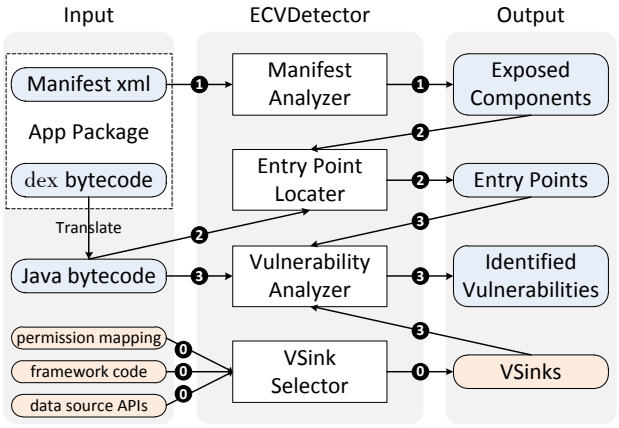

Figure 5: ECVDetector architecture. VSink Selector and Vulnerability Analyzer are two major components to implement our sink-driven approach.

Among them, 137 APIs are `VS_Direct`, 23 `VS_DirectByParam`, 167 `VS_Input`, and 45 `VS_Public`. Moreover, to facilitate detecting ECVs involving sensitive sources, we also select 183 `VSource` based on the results of Stowaway and SuSi [2]. As regards to the data sinks in SuSi, we find most of them cannot serve our vulnerability-specific purpose, due to their selection motivation for privacy leak detection.

To adopt the Stowaway mapping into our VSinks, however, faces a technical challenge that method signatures in Stowaway mapping are incomplete. In fact, these incomplete signatures are only sub-signatures without return-value types, thus would incur inaccuracy for Vulnerability Analyzer. To this end, we design several estimation rules to restore complete method signatures by analyzing corresponding Android framework code. Moreover, since our analysis program requires right method sub-signatures to facilitate signature restoring, we unexpectedly identify several kinds of inaccuracy issues existing in Stowaway results. These issues exist might because Stowaway also employs some manual efforts to produce the mapping, thus introduce some inaccuracies. We then ran our analysis program several times to inspect the error messages (in each run), which prevents us to successfully restore all signatures. For identified issues in each run, we manually fix them by batch replacing.

Manifest Analyzer. This component consists of two parts: an xml parser to extract all essential information inside each manifest file, and a script to identify exposed components by analyzing extracted information. We determine exposed components according to the rule in §2.1.

Entry Point Locator. The main task of Entry Point Locator is to locate all entry points, which would serve as the starting points for Vulnerability Analyzer. Basically, entry points are fixed callback interfaces defined by Android programming paradigm, such as onCreate and onStart. Particularly, onCreate is an initialization point which will be called when a component is started up. Moreover, several kinds of entry points are special, in terms of they can accept external attack inputs. While other points either take zero parameter (e.g., onResume and onStop), or only contain inputs cannot be manipulated by attacker (e.g., `Bundle` argument in Activity's onCreate entry point). Since `VS_Input` ECV detection is related to those special entry points, we identify them in Android framework and list in Table 1. Three kinds of components contain such entry points we concerned, while Activity needs to actively call `getIntent` fucntion to receive external inputs.

Table 1: Our aimed entry point functions.

Component	Callback functions that accept attack inputs
Service	onBind(**Intent** intent) onStart(**Intent** intent, int startId) onStartCommand(**Intent** intent, int flags, int Id) onHandleIntent(**Intent** intent) handleMessage(**Message** msg)
Receiver	onReceive(Context context, **Intent** intent)
Provider	query(**Uri**, String[], String, String[], String) insert(**Uri**, ContentValues) update(**Uri**, ContentValues, String, String[]) delete(**Uri**, String, String[]) openFile (**Uri**, String)

Another task is to model lifecycle for identified entry points. This is necessary because the entry points will not call each other due to their callback nature, thus cause disconnected static flows. Moreover, there are some initialization functions even before the onCreate interface, such as <clinit> and <init>. We also need to consider them in modeling lifecycle, so that backward module could backtrack to the initial definition. In the current ECVDetector, we model the lifecycle by defining several continuous phases. Each phase may contain a manually connected flow or several asynchronous entry points. For example, we define an "initial" phase to connect the initial flow (i.e., <clinit> → <init> → onCreate), and a "main" phase to cover all asynchronous special entry points in Table 1.

Vulnerability Analyzer. This component mainly performs the sink-driven forward and backward analysis in §3.2. In the process of implementing Vulnerability Analyzer, we come across two major technical issues due to: (1) object-oriented (OO) language used by Android development, and (2) Android's event-driven nature.

The first issue arises from the inheritance nature of OO language during the call site resolving in the forward module. To obtain the target method of a call site, it is essential to resolve the type of target class object. However, due to the inheritance, it is hard to statically determine what concrete class an object would represent. More specifically, the inheritance nature allows developer to use superclass or interface type to represent a subclass object, which causes the ambiguity. We tackle this problem to a great extent by leveraging typed Soot IR [24], which is calculated by a fast type

inference algorithm [4]. Besides the object type resolving, OO language's inheritance nature also makes locating the right method definition complex, even if we have obtained the exact object type. For example, `toString` method invoked by `android.graphics.Bitmap` object is actually not defined by `Bitmap` class itself. Indeed, we have to track back to `java.lang.Object` class to obtain the method definition of `toString`. A straightforward way for mitigating this problem is to maintain a comprehensive class hierarchy. In our ECVDetector prototype, we also take advantage of the automatic method resolution provided by Soot according to Java Virtual Machine Specification [25].

The second issue is the static flow discontinuity problem caused by Android's event-driven nature [19]. A typical disconnected example is the flow between `start()` and `run()` method in `java.lang.Thread` class. In fact, Android OS connects these two methods dynamically through the underlying thread scheduler. However, a static tool, without specific knowledge, cannot directly predict this kind of dynamic connecting behavior. Nevertheless, it is necessary for Android static analysis tools to tackle this problem in a satisfactory way. Otherwise, some false negatives would arise. In ECVDetector, we model a number of dynamic flow connecting behaviors as pre-defined knowledge to build continuous call chains. The main modeled flow connecting behaviors are summarized in Table 2. According to this table, ECVDetector is capable to connect the disconnected flows occurred in thread scheduling, timers, location updates, and so on.

Table 2: The main dynamic flow connecting behaviors modeled by ECVDetector.

Class name	Modeled Flows
Thread	start → run
AsyncTask	execute → doInBackground
Handler	sendMessage → handleMessage
Timer	schedule → run
LocationManager	requestLocationUpdates → *LocationListener*
TelephonyManager	listen → *PhoneStateListener*

5. EVALUATION

To evaluate the efficacy and performance of ECVDetector, we carried out an evaluation with top 1K Android apps from Google Play. The reason we evaluated these popular apps is that the impact of ECVs relies on the popularity of vulnerable apps. In other words, vulnerable apps with few installs would not cause big security impact, because it is less likely to let those limited users also install the attack apps. Specifically, these top apps were selected according to their user review numbers, and most of them were crawled recently (between June and July 2013). Therefore, our app dataset could represent the recent versions of top Google Play apps that users may install in their phones. Moreover, our selected apps were fully unique in terms of the package names, which made our dataset more distributed.

In this section, we first discuss how we conducted the experiment with this dataset, including some findings and knowledge we obtained during the procedure of experiment. Then we report our identified ECVs and conduct case studies of some representative vulnerable apps. Finally, we depict the result of performance evaluation for ECVDetector.

5.1 Experiment and Findings

One essential step before the ECVDetector experiment is to extract the AndroidManifest.xml for each app. As mentioned in §4, we took apktool for unpacking the app and obtaining the manifest file. However, we found not all apps could be correctly handled by apktool. Indeed, six apps of our dataset fail and crash apktool with several Java exceptions, maybe due to the potential bugs in apktool or some apps try to protect itself from the decompiling of apktool [30]. This finding also gives a kind hint to all previous work based on apktool, such as [12] and [20].

We then leveraged the Manifest Analyzer component of ECVDetector to discover exposed components from the rest of 994 apps. Among them, we successfully parsed and analyzed 992 manifest files. The two failure cases were due to the invalid encodings that cannot be handled by the Python XML DOM library. More specifically, one failed case (the popular Titanium Backup app) contains the Chinese and Korean characters for some component names. These two failed manifest files could be manually analyzed for obtaining exposed components, but in this paper we just ignored them for the automatic experiment.

In summary, we identified a total of 7,664 exposed components from the remaining 992 apps, and 6,582 of them were unique in terms of the component name. The detailed amounts of exposed components classified by component types are illustrated in Figure 6. One major finding based on this figure is that there are significantly more Activity and Broadcast Receiver exposed components than the other two component types, around ten times. The reason behind this can be explained by these two facts. First, many apps need exposed Activities to finish the user intention of app switching, such as launching from the launcher app. Second, in order to receive system broadcasts, the corresponding Receivers have to expose themselves to the Android framework.

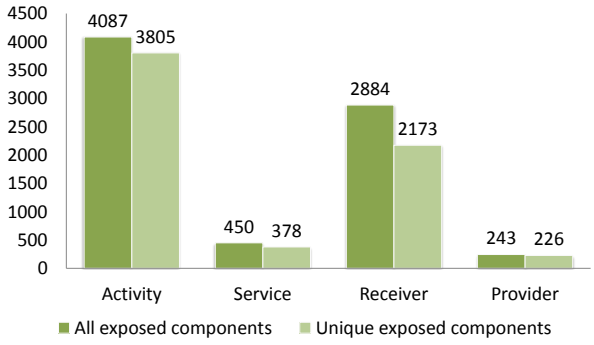

Figure 6: The amounts of all and unique exposed components across four component types.

Our experiment for ECV detection focused on the exposed Services and Broadcast Receivers. More specifically, our target was 2,551 unique exposed components, including 378 Services and 2,173 Receivers (see Figure 6). The reason for skipping Activities is mainly because exploiting Activity vulnerabilities usually requires user's intention (e.g., the Adobe Activity vulnerability shown in [7]), and the complicated attacks against Browser Activities [31] are out of the scope of our paper (see §2.2). As for Content Providers, although ECVDetector is capable of statically detecting some potential Provider vulnerabilities, heavy manual efforts are still

required even with the dedicated ContentScope [36] tool.

Next, we ran ECVDetector against those 2,551 exposed components. The experiment was performed on a single Dell PowerEdge storage server, equipped with four 2.40GHz CPUs and 12GB of RAM. To optimize the throughout, we set the timeout for processing each component to two minutes. This threshold was quite reasonable, because we found most of components could be analyzed within 10 seconds (see §5.3 later).

During the experiment, we found it was challenging to prevent Soot from crashing when loading Java classes for analysis. We spent many efforts to make ECVDetector successfully analyze all 2,551 components, and our knowledge listed as follows might be helpful other Soot-based Android tools [16, 28, 34]. The first knowledge is that we should provide as many underlying Android classes as possible to Soot. These classes would help Soot resolve most of Android-specific types. Specifically, we provided eight platforms of Android SDKs to Soot, from the old 1.6 to the recent 4.3. Moreover, two platforms of hidden and system Android APIs were also provided to Soot, as well as one version of the Google APIs (e.g., Google Map APIs). However, we still encountered many failed cases due to missing app-specific classes. In order to solve this issue, we made use of the second knowledge: ask Soot to only load classes on demand, and if errors still occur then set the Soot option `allow_phantom_refs`. This option would let Soot create a phantom class for each missing class.

Another interesting finding is about system-only broadcasts, which are discussed in §3.2.3. Our finding about these broadcasts for the 2173 exposed Receivers is divided into three parts. First, we identify 433 of them contained the checking for system-only broadcasts in their codes. This result suggests our designed validation capability (§3.2.3) could effectively alert the false positives due to system-only broadcast checking. Second, total 40 broadcasts of our selected 133 system-only broadcasts are checked. Third, we also count the numbers of checked broadcasts, and the top three result is listed in Table 3. The fourth most checked broadcast is `TIMEZONE_CHANGED`, but with only 19 times.

Table 3: Top 3 checked system-only broadcasts.

The Checked System-only Broadcast	#
android.intent.action.BOOT_COMPLETED	157
android.net.conn.CONNECTIVITY_CHANGE	57
android.intent.action.PACKAGE_ADDED	43

Through the automatic analysis for the 2,551 exposed components of 992 apps, we discovered the initial set of potential ECVs with 348 affected apps. As previously shown in Figure 1, this result was analyzed only by the forward and backward analysis in ECVDetector. We still needed to perform the semi-auto guided analysis for two ECV categories (`VS_DirectByParam` and `VS_Input`), as discussed in §3.3. Specifically, there were 172 and 357 unique logs (need guided analysis) for `VS_DirectByParam` and `VS_Input`, respectively. Nevertheless, most of them were similar to each other, only 6 kinds of `VS_DirectByParam` APIs and 16 kinds of `VS_Input` APIs were identified. Therefore, we did not spend much effort (around one hour in total) in validating these logs. ECVDetector then automatically filtered various sink-specific false positives according to our extracted log pattern. Eventually, we identified a total of 103 potential vulnerable apps. More specifically, there are 31 apps affected with potential `VS_Direct` ECVs, 25 for `VS_DirectByParam`, 56 for `VS_Input`, and 16 for `VS_Public`.

5.2 Identified ECVs

The 103 potential vulnerable apps (identified in the last section) were further manually verified for the real ECVs, due to the lack of ground truth. The result of this manual verification mainly relies on the domain knowledge of security analysts and their defined boundary about bug and vulnerability, as discussed in §2.2. In our verification process, we tried to follow a relatively more conservative principle than the closest related work CHEX [26]. A typical example is that the vulnerable case of "Object embedded in input used to start Activity" in CHEX would not be directly treated as a vulnerability in our principle. Instead, we also require this action of starting Activity could make some security impact to the internal status of victim apps for becoming an ECV. Some related preliminary knowledge and discussion could be also found in §2.2.

Overall results. In the end, we tagged 49 apps as vulnerable from 103 candidate ones. We confirmed them mainly by carefully auditing the decompiled codes and manifest files (with the help of result logs and call chains produced by ECVDetector), similar to the verification used by CHEX. Note that for each vulnerable app, we only tagged it to one major ECV category, even if this app might be affected by several ECVs. The main result is summarized in Table 4, including the number of each category of identified ECVs and their corresponding representative vulnerable apps. To the best of our knowledge, most of them are not disclosed previously, thus are the zero-day ECVs. Since these vulnerable apps are from top 1K in Google Play, their vulnerabilities may incur real security consequences among many users. We have reported several particularly serious cases to corresponding vendors, and received the acknowledgements from Go Dev team and LinkedIn.

The true positive rate of ECVDetector is not high (i.e., 47.6%, 49/103), but we argue this is acceptable when consider the following two factors. The first factor is the aforementioned conservative verification principle we used, which cause us to report less vulnerabilities. Some `VS_Input` APIs are the most affected by this principle, mainly the IPC APIs (e.g., `startActivity` and `startService`). This problem could be migrated when further cross-component inspection is involved. The second is due to the characteristics of some our selected sinks. For instance, the `removeAccount` `VS_Direct` API is commonly used in a reasonable scenario—when an user logs out her account, the corresponding app then removes her local account cache from the phone—which should not be treated as an vulnerability. However, it is hard for ECVDetector to recognize this normal situation of API usage.

Case studies. We now conduct three cases studies to demonstrate ECVDetector's capability of detecting real serious ECVs under different categories. For each case, we further write an exploit in a zero-permission app to confirm the corresponding ECV could be effectively exploited. Automatically generating exploits is an interesting topic for aiding static analysis, but is out of our paper scope and will be a future work. Here we hide app version information in case that some users are still using the vulnerable versions

Table 4: Four categories of identified ECVs and their representative vulnerable apps.

ECV Category	# of Apps	Representative Apps	
		Package Name	Vulnerability Description
VS_Direct	6	com.jb.gosms	Force it to send SMS to phone no. specified by input
		com.gau.go.launcherex.gowidget.switchwidget	Force it to enable or disable wifi and bluetooth
		com.bwx.bequick	Force it to open the camera as flash light
VS_DirectByParam	5	com.cleanmaster.mguard	Force it to silently uninstall arbitrary apps and etc.
		mominis.Generic_Android.Ninja_Chicken	Force it to delete its internal "MESSAGE" database table
VS_Input	25	com.zlango.zms	Force it to change the status of messages in SMS databases
		com.antivirus	Start the private core AVService with extras specified by input
		com.doubleTwist.androidPlayer	Start a private service with dangerous action named "delete_db"
		com.linkedin.android	Launch a network request with the attributes specified by input
		com.ebuddy.android	The beta update Receiver can be cheated to download fake apps
VS_Public	13	com.symantec.mobilesecurity	getLastKnownLocation() is outputted to the logcat
		air.com.bitrhymes.bingo	getDeviceId() along with post URL are outputted to the logcat
		com.levelup.touiteur	getAllNetworkInfo() is outputted to the logcat
		com.sec.spp.push	getConnectionInfo() is outputted to the logcat

of apps.

Go SMS Pro (`com.jb.gosms`) is the top 1 messaging app with over 60 million installs. ECVDetector identified its exposed `CellValidateService` component leaked a security-critical flow path which arrived at the `sendTextMessage` sink API under the `VS_Direct` category. Moreover, its first parameter for assigning target phone number is completely controllable by external Intent. Thus, a zero-permission attack app can force Go SMS Pro to send SMS to arbitrary phone number. We also prepared a demo video (at `https://www.youtube.com/watch?v=CwtNCwAHSRs`) and reported it to Go Dev team on Sep 9 2013. From their active response, we knew we were the first reporter on this issue. This demonstrates the efficacy of ECVDetector, even when `sendTextMessage` is a commonly-used sink.

Clean Master (`com.cleanmaster.mguard`) is a quite popular clean up app with over 200 million worldwide installs. ECVDetector identified an external Intent can control its exposed `LocalService` to do privileged operations, such as clean up memory, restore apps, and silently uninstall arbitrary apps. Take silently uninstalling apps as an example, the flow to a `VS_DirectByParam` command execution API could be injected into attacker-supplied parameters (e.g., a victim app). Clean Master then executes the "pm uninstall" command to do silent uninstallation (of course, root permission needs to be pre-granted). This case demonstrates ECVDetector can uncover serious `VS_DirectByParam` vulnerabilities.

Lango Messaging (`com.zlango.zms`) is another top messaging app with over one million installs. ECVDetector identified an external Intent (at its `data` field) can flow into a `VS_Input` API (called `ContentResolver.update`) via the exposed `ZmsSentReceiver` component. We further found this exposed sensitive flow can enable a zero-permission app to maliciously change the status of some SMS messages. For example, a draft SMS can be changed into the "sent" status and this changing will be reflected in the UI of all installed messaging apps. Even worse, an incoming SMS can be set as an outgoing message, thus indirectly producing fake SMS messages. This case shows the capability of ECVDetector to detect `VS_Input` ECVs.

However, generating an effective exploit for Lango Messaging is surprisingly challenging, nearly taking us one-day effort. The major difficulty is how to obtain an appropriate `Uri` data field from infinite candidates with the format of "`content://`". The only hint is a domain knowledge required condition judgment (`MessageItem.getBoxId() != 2`). Eventually, we figure out the `MessageItem` object (from `MessageItemManager.get(Uri)`) refers to a SMS message, and its status should be not as "sent". Therefore, a target `Uri` can be "`content://sms/inbox||draft/id`". The next step is to calculate a right message id. Since the attack app has no `READ_SMS` permission, it needs to launch brute-force attempts. In summary, this case shows static detection tools (like ECVDetector) and manual analysis are still necessary. By contrast, automatic exploit generation (like fuzzing [27, 33] or symbolic execution [35]) is nearly impossible for handling this case, due to the lack of domain knowledge and setting up related system conditions (in this case, some incoming and draft SMS have to be prepared).

`VS_Public` ECVs are relatively not so serious, and require more app-specific conditions to trigger the vulnerable paths. Therefore, we skip their case studies in this paper.

5.3 Performance Evaluation

ECVDetector spent 6h 33m 35s to analyze the total 2,551 exposed components. Therefore, the average processing time for each component was 9.257s. This result suggests the threshold time of two minutes is quite reasonable. Among the 2,551 components, we found only 35 of them arrived this timeout value. The reason might be the complicated recursion codes that ECVDetector cannot recognize, although we have designed some mechanisms to help ECVDetector handle simple recursion. We further give the detailed performance measurement for the rest of 2,516 components, as shown in Figure 7. It is easy to determine that most of com-

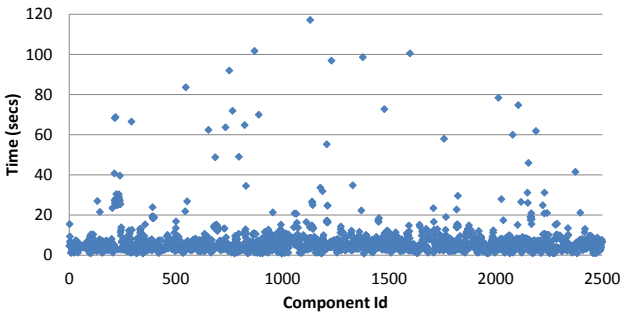

Figure 7: Detailed performance measurement.

ponents could be analyzed within 10 seconds. Therefore, we conclude that the performance of ECVDetector is high.

6. DISCUSSION

The false positives shown in the last section are mainly due to the variable semantics of the `VS_DirectByParam` and `VS_Input` APIs. It is hard to reduce these false positives even with dynamic analysis, because they also cannot determine which parameter values are critical and harmful. However, security analysts can extract the pattern of false positives from existing manual analysis, and feed them into the automatic filtering module of our guided analysis (§3.3). In this way, some false positives could be further excluded automatically.

Similar to other works, there are also some false negatives in the current ECVDetector prototype. However, it is hard to measure their quantity due to the lack of ground truth. Thus, we discuss several possible causes for these false negatives as follows. First, our modeling for connecting dynamic flows (see Table 2) is not intended to be complete, but we have tried our best to cover the common cases, similar to [19, 36]. Second, although we have proposed a systematic strategy to collect many sinks, it is impossible to cover a complete set. Third, our current implementation of backtracking call chains is also not prefect in that we skip analyzing other callee methods in the same level of hierarchy for code simplicity and performance consideration. Fortunately, only a small number produce false negatives is due to this limitation. Finally, the cross-component and cross-app ECVs [32] are not considered in our work, but we could leverage Epicc [28] to build a knowledge database and handle these cases.

7. RELATED WORK

Exposed Component Vulnerability. This class of vulnerabilities has attracted much research effort recently, after Davi et al. [9] presents the basic ECV model focusing on permission leaks. Subsequently, several detection methods with different foci have been proposed. One category of these methods includes ComDroid [8] and Epicc [28], both of them try to identify several potential security flaws during the Intent-based communication. As mentioned in §2.2, they aim at discovering all attack surfaces and give security warnings. Therefore, they do not conduct sink-based analysis as ours. As a result, they issue many potential ECV warnings, but it is hard to uncover true ECVs that have security impacts.

Another category of methods performs sink-based flow analysis to identify vulnerabilities more accurately. Specifically, Woodpecker [19] and DroidChecker [7] leverage path reachability analysis to detect capability or permission leaks, which generally belong to our `VS_Direct` ECV category except for the non-ECV vulnerabilities on unauthorized Intent receipt. SEFA [32] further extends this idea by adopting content leak detection [36] in exposed Content Providers. CHEX [26], on the other hand, employs inter-procedural and context-sensitive dataflow techniques to locate suspicious ECV flows terminating at the data sinks (i.e., our `VS_Input` and `VS_Public` APIs).

Similar to other sink-based analysis systems, ECVDetector also conducts path reachability and dataflow analysis. However, unlike previous works, these two kinds of analysis techniques are organically combined together and driven by the classified categories of VSinks in our approach. In particular, we take the backward, instead of previous forward, dataflow analysis for adapting more categories of sinks, such as the `VS_DirectByParam` category. These categorized VSinks are obtained by our systematic VSink selection strategy, which also assists ECVDetector to cover more ECVs that are not addressed by previous work. Additionally, we further design semi-auto guided analysis and system-only broadcast checking capability to efficiently exclude some false positives.

Besides detection techniques, several defense methods are proposed to mitigate ECV issues. In general, they are dynamic enforcing systems, therefore required to modify Android source code. They propose to check IPC call chains [15, 10], or even other channels like sockets and files [5]. On the other hand, Kantola et al. [23] tries to automatically reduce unnecessary exposed surfaces. Moreover, mandatory access control is tailored to Android [29, 6], and it could block ECV attacks with appropriate policies. Quite recently, AppSealer [34] aims to automatically generate patches for preventing attacks to exploit ECVs. All these works are complementary to our work, since they might be applied to prevent attackers from exploiting our discovered ECVs.

Sink Selection in App Analysis. Many app analysis tools rely on sinks to perform their individual analysis. However, most of their sinks are selected in a manual fashion [13, 22, 11, 19, 36]. Some works make use of API-to-permission mappings [14, 3] to obtain their target sinks, such as [17, 18, 26]. However, due to the lack of systematic strategy and flexible rules like our own, it is not easy for them to systematically extract appropriate sinks and filter useless ones from all candidate mappings. Moreover, we craft centralized rules to adopt the privileged APIs not covered by existing mappings, as well as other non-privileged APIs like database APIs. Another related work is SuSi [2], which employed machine learning techniques to select data sinks for privacy leak detection. However, we cannot use SuSi to select our VSinks because of the variable API semantics and that we also consider non-data sinks.

A contribution of this paper is the semi-auto VSink classification. SuSi and CHEX also define categories for their selected sinks, but their categories are not meant for capturing different analysis requirements like ours. In fact, the categories in SuSi are only the API types (e.g., Network and File sinks). Similarly, CHEX defines three sink tags for just differentiating different data sinks. On the other hand, although both DroidChecker [7] and AdRisk [18] separate their sinks into two categories for analysis, they are not

fully aimed at detecting ECV detection, thus not sufficient for all analysis requirements in ECV detection. Indeed, one of their categories is either for detecting unauthorized Intent receipt or for privacy leak detection.

8. CONCLUSION

In this paper, we presented a new sink-driven approach to systematically tackle the ECV detection problem. This approach includes a systematic strategy for VSink selection and classification, and a general detection method to identify potential ECVs in Android apps. We implemented our sink-driven approach in a tool called ECVDetector. We successfully identified a total of 49 vulnerable apps across all four ECV categories in the top 1K Android apps. Future works include helping developers fix their vulnerable apps and deploying ECVDetector as a web-based detection service.

9. REFERENCES

[1] Android Application Components. `http://developer.android.com/guide/components/fundamentals.html#Components`.

[2] S. Arzt, S. Rasthofer, and E. Bodden. Susi: A tool for the fully automated classification and categorization of Android sources and sinks. In *Technical Report TUD-CS-2013-0114*, 2013.

[3] K. Au, Y. Zhou, Z. Huang, and D. Lie. Pscout: Analyzing the Android permission specification. In *Proc. ACM CCS*, 2012.

[4] B. Bellamy, P. Avgustinov, O. Moor, and D. Sereni. Efficient local type inference. In *Proc. ACM OOPSLA*, 2008.

[5] S. Bugiel, L. Davi, A. Dmitrienko, T. Fischer, A. Sadeghi, and B. Shastry. Towards taming privilege-escalation attacks on Android. In *Proc. ISOC NDSS*, 2012.

[6] S. Bugiel, S. Heuser, and A. Sadeghi. Flexible and fine-grained mandatory access control on Android for diverse security and privacy policies. In *Proc. Usenix Security*, 2013.

[7] P. Chan, L. Hui, and S. Yiu. Droidchecker: Analyzing Android applications for capability leak. In *Proc. ACM WiSec*, 2012.

[8] E. Chin, A. P. Felt, K. Greenwood, and D. Wagner. Analyzing inter-application communication in Android. In *Proc. ACM MobiSys*, 2011.

[9] L. Davi, A. Dmitrienko, A. Sadeghi, and M. Winandy. Privilege escalation attacks on Android. In *Proc. Springer ISC*, 2010.

[10] M. Dietz, S. Shekhar, Y. Pisetsky, A. Shu, and D. Wallach. Quire: Lightweight provenance for smart phone operating systems. In *Proc. USENIX Security*, 2011.

[11] M. Egele, C. Kruegel, E. Kirda, and G. Vigna. Pios: Detecting privacy leaks in iOS applications. In *Proc. ISOC NDSS*, 2011.

[12] K. Elish, D. Yao, B. Ryder, and X. Jiang. A static assurance analysis of Android applications. `http://people.cs.vt.edu/~danfeng/papers/user-intention-PA-2013.pdf`, 2013.

[13] W. Enck, P. Gilbert, B. Chun, L. Cox, J. Jung, P. McDaniel, and A. Sheth. Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. In *Proc. Usenix OSDI*, 2010.

[14] A. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *Proc. ACM CCS*, 2011.

[15] A. Felt, H. Wang, A. Moshchuk, S. Hanna, and E. Chin. Permission re-delegation: Attacks and defenses. In *Proc. Usenix Security*, 2011.

[16] C. Fritz, S. Arzt, S. Rasthofer, E. Bodden, A. Bartel, J. Klein, Y. Traon, D. Octeau, and P. McDaniel. Highly precise taint analysis for Android application. In *Technical Report TUD-CS-2013-0113*, 2013.

[17] C. Gibler, J. Crussell, J. Erickson, and H. Chen. Androidleaks: Automatically detecting potential privacy leaks in Android applications on a large scale. In *Proc. Springer TRUST*, 2012.

[18] M. Grace, W. Zhou, X. Jiang, and A. Sadeghi. Unsafe exposure analysis of mobile in-app advertisements. In *Proc. ACM WiSec*, 2012.

[19] M. Grace, Y. Zhou, Z. Wang, and X. Jiang. Systematic detection of capability leaks in stock Android smartphones. In *Proc. ISOC NDSS*, 2012.

[20] M. Grace, Y. Zhou, Q. Zhang, S. Zou, and X. Jiang. Riskranker: Scalable and accurate zero-day Android malware detection. In *Proc. ACM MobiSys*, 2012.

[21] N. Hardy. The confused deputy: (or why capabilities might have been invented). In *ACM SIGPOS Operating Systems Review*, 1988.

[22] P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall. "these aren't the droids you're looking for": Retroffiting Android to protect data from imperious applications. In *Proc. ACM CCS*, 2011.

[23] D. Kantola, E. Chin, W. He, and D. Wagner. Reducing attack surfaces for intra-application communication in Android. In *Proc. ACM SPSM*, 2012.

[24] P. Lam, E. Bodden, O. Lhotak, and L. Hendren. The soot framework for java program analysis: a retrospective. In *Cetus Users and Compiler Infastructure Workshop (CETUS 2011)*, 2011.

[25] T. Lindholm, F. Yellin, G. Bracha, and A. Buckley. The java virtual machine specification, java se 7 edition. `http://docs.oracle.com/javase/specs/jvms/se7/jvms7.pdf`, 2012.

[26] L. Lu, Z. Li, Z. Wu, W. Lee, and G. Jiang. Chex: Statically vetting Android apps for component hijacking vulnerabilities. In *Proc. ACM CCS*, 2012.

[27] A. Maji, F. Arshad, S. Bagchi, and J. Rellermeyer. An empirical study of the robustness of inter-component communication in Android. In *Proc. IEEE DSN*, 2012.

[28] D. Octeau, P. McDaniel, S. Jha, A. Bartel, E. Bodden, J. Klein, and Y. Traon. Effective inter-component communication mapping in Android with Epicc: An essential step towards holistic security analysis. In *Proc. Usenix Security*, 2013.

[29] S. Smalley and R. Craig. Security enhanced (se) Android: Bringing flexible mac to android. In *Proc. ISOC NDSS*, 2013.

[30] T. Strazzere. Slowing down Android reverse engineers. `http://files.meetup.com/3970892/SlowingReversers.pdf`, 2013.

[31] R. Wang, L. Xing, X. Wang, and S. Chen. Unauthorized origin crossing on mobile platforms: Threats and mitigation. In *Proc. ACM CCS*, 2013.

[32] L. Wu, M. Grace, Y. Zhou, C. Wu, and X. Jiang. The impact of vendor customizations on Android security. In *Proc. ACM CCS*, 2013.

[33] K. Yang, L. Zhou, Y. Wang, J. Zhuge, and H. Duan. Intentfuzzer: Detecting capability leaks of Android applications. In *Proc. ACM AsiaCCS*, 2014.

[34] M. Zhang and H. Yin. Appsealer: Automatic generation of vulnerability-specific patches for preventing component hijacking attacks in Android applications. In *Proc. ISOC NDSS*, 2014.

[35] J. Zhong, J. Huang, and B. Liang. Android permission re-delegation detection and test case generation. In *Proc. IEEE International Conference on Computer Science and Service System*, 2012.

[36] Y. Zhou and X. Jiang. Detecting passive content leaks and pollution in Android applications. In *Proc. ISOC NDSS*, 2013.

APPENDIX

A. SYSTEM-ONLY BROADCAST CHECKING

```
<receiver android:name="com.example.BootReceiver">
   <intent-filter>
      <action android:name=
            "android.intent.action.BOOT_COMPLETED"/>
   </intent-filter>
</receiver>
```

Figure 8: An example Broadcast Receiver registering the `BOOT_COMPLETED` broadcast.

```
1 public class BootReceiver extends BroadcastReceiver {
2
3   public void onReceive(Context context, Intent intent) {
4     if ("android.intent.action.BOOT_COMPLETED".
          equals(intent.getAction())) {
5         // other codes
6     }
7   }
8 }
```

Figure 9: A typical code pattern for checking system-only broadcasts.